

Intel[®] True Scale Fabric Suite Software

Release Notes

February 2018



No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors which may cause deviations from published specifications.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting: <http://www.intel.com/design/literature.htm>

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.


*Other names and brands may be claimed as the property of others.

Copyright © 2018, Intel Corporation. All rights reserved.




Contents

1.0	Overview of the Release	7
1.1	Introduction	7
1.2	Audience	7
1.3	If You Need Help	7
1.4	New Features and Enhancements	7
1.4.1	Release 7.4.4.0.18 Enhancements	7
1.4.2	Release 7.4.2.0.6 Enhancements	8
1.5	Operating Environments Supported	8
1.6	Qualified Parallel File Systems	9
1.7	Intel Interface for NVIDIA* GPUs	9
1.8	Compilers Supported	9
1.8.1	MPI	9
1.8.2	MVAPICH, MVAPICH2 and Open MPI	10
1.9	Hardware Supported	10
1.10	Software Supported	11
1.10.1	Remote Node Software Versions Supported in this Release	11
1.10.2	Remote Node Software Versions with Reduced Capability	11
1.11	Installation Requirements	11
1.11.1	Package Installation Requirements:	12
1.11.2	Software and Firmware Requirements	12
1.12	Changes for this Release	12
1.12.1	Changes to Industry Standards Compliance	12
1.13	Product Constraints	12
1.13.1	FastFabric Toolset Product Constraints	12
1.13.2	Fabric Manager	13
1.14	Product Limitations	13
1.15	Other Information	13
1.15.1	FastFabric Toolset Information	13
1.15.2	Fabric Manager Information	14
1.15.3	Verbs Performance	14
1.16	Documentation	14
2.0	System Issues for Release 7.4.4.0.18	15
2.1	Introduction	15
2.2	Resolved Issues in this Release	15
2.3	Known Issues	16



Tables

1-1 Operating Environments Supported	8
1-2 CPU Model of Linux Kernel	8
1-3 NVIDIA's CUDA Tested with IFS	9
1-4 MPI Compilers	9
1-5 MVAPICH, MVAPICH2 and Open MPI	10
1-6 Hardware Supported	10
1-7 Changes to Industry Standards Compliance	12
1-8 Related Documentation for this Release	14
2-1 Resolved Issues	15
2-2 Open Issues	16



1.0 Overview of the Release

1.1 Introduction

These Release Notes provide a brief overview of the changes introduced into the Intel® True Scale Fabric Suite Software (IFS) software by this release. This release notes document includes only the IFS software and must be used in conjunction with the Intel® *OFED+ Host Software Release Notes* for a complete package. References to more detailed information are provided where necessary. The information contained in this document is intended for supplemental use only; it should be used in conjunction with the documentation provided for each component.

These Release Notes list the new features of the release, as well as the system issues that were closed in the development of Release 7.4.4.0.18.

1.2 Audience

The information provided in this document is intended for installers, software support engineers, and service personnel.

1.3 If You Need Help

If you need assistance while working with the True Scale Fabric Suite Software, contact your Intel approved reseller or Intel® True Scale Technical Support:

- By E-mail:
fabricsupport@intel.com
- On the Support tab at website:
[Fabric Products](#)

For OEM-specific server platforms supported by this release, contact your OEM.

1.4 New Features and Enhancements

This section list the new features and enhancements for this release, as well as the previous release.

1.4.1 Release 7.4.4.0.18 Enhancements

- Added support for
 - RHEL 7.3
 - RHEL 7.4
 - SLES 12.2
 - CentOS 7.3
 - CentOS 7.4



- Scientific Linux 7.3
- Scientific Linux 7.4
- OFED 4.8-1

1.4.2 Release 7.4.2.0.6 Enhancements

- Added support for
 - RHEL 6.8
 - CentOS 6.8
 - Scientific Linux 6.8
 - OFED-3.18-2

1.5 Operating Environments Supported

The following Operating Systems are supported in this release:

Table 1-1. Operating Environments Supported

Operating System	Update/SP	Version
RHEL 7 X86_64 (AMD Opteron and Intel EM64T)	Update 3	3.10.0-514.el7.x86_64
	Update 4	3.10.0-693.el7.x86_64
SLES 12 X86_64 (AMD Opteron and Intel EM64T)	Update 1	3.12.49-11.1-default
	Update 2	4.4.21-69-default
Community Enterprise Operating System (CentOS) X86_64 (AMD Opteron and Intel EM64T) (7.x)	Update 3	3.10.0-514.el7.x86_64
	Update 4	3.10.0-693.el7.x86_64
Scientific Linux X86_64 (7.x)	Update 3	3.10.0-514.el7.x86_64
	Update 4	3.10.0-693.el7.x86_64

CPU model of Linux kernel can be identified by `uname -m` and `/proc/cpuinfo` shown in [Table 1-2](#).

Table 1-2. CPU Model of Linux Kernel

Model	uname	/proc/cpuinfo
EM64T	x86_64	Intel CPUs
Opteron*	x86_64	AMD CPUs

Note: Other combinations (such as i586 `uname`) are not currently supported.



1.6 Qualified Parallel File Systems

Lustre and IBM* General Parallel File System (GPFS) listed below have been tested for use with this release of the Intel® OFED+ host software using the operating systems listed below:

- Lustre* 2.10.1
 - RHEL 7.4 (both server and client)
 - SLES12.2 (client)
- IBM GPFS 4.2.3
 - RHEL 7.2
 - RHEL 7.3
 - SLES 12.1
 - SLES 12.2

Refer to the *Intel® OFED+ Host Software User Guide* for the latest configuration recommendations for optimizing Lustre and GPFS performance with Intel® True Scale Fabric.

1.7 Intel Interface for NVIDIA* GPUs

NVIDIA's CUDA parallel computing platform and programming models have been tested for use with this release of the Intel® OFED+ host software using the operating systems listed in [Table 1-3](#):

Table 1-3. NVIDIA's CUDA Tested with IFS

Distributions	CUDA 7
RHEL 7.2	X

1.8 Compilers Supported

1.8.1 MPI

This release supports the following MPI implementations:

Table 1-4. MPI Compilers

MPI Implementation	Runs Over	Compiled With
Open MPI 1.10.7	PSM Verbs	GCC, Intel, PGI
MVAPICH2 version 2.1.1	PSM Verbs	GCC, Intel, PGI
Platform MPI 9.1	PSM Verbs	GCC (default)
Intel MPI version 2017 Build 2017.4.196	TMI/PSM, uDAPL	icc 17.0.4.196



1.8.2 MVAPICH2 and Open MPI

MVAPICH2 and Open MPI have been compiled for PSM to support the following versions of the compilers:

Table 1-5. MVAPICH2 and Open MPI

Compiler name	Distro	Compiler Version
(GNU) gcc	RHEL6	gcc (GCC) 4.4.7 20120313 (Red Hat 4.4.7-17)
(GNU) gcc	RHEL7	gcc (GCC) 4.8.5 20150623 (Red Hat 4.8.5-4)
(GNU) gcc	SLES11	gcc (SUSE Linux) 4.3.4 [gcc-4_3-branch revision 152973]
(PGI) pgcc	RHEL6	pgcc 2017-174 64-bit target on x86-64 Linux -tp nehalem-64
(PGI) pgcc	RHEL7	pgcc 2017-174 64-bit target on x86-64 Linux -tp sandybridge
(PGI) pgcc	SLES12	pgcc 2017-174 64-bit target on x86-64 Linux -tp nehalem-64
(Intel) icc	RHEL7	17.0.4.196
(Intel) icc	SLES12.1	17.0.4.196

1.9 Hardware Supported

Table 1-6 list the hardware supported in this release.

Table 1-6. Hardware Supported

HCA
QLE7340
QLE7342
QME7342
QME7362
QMH7342
MHQH29-*
MHQH19-*
MHQH19B-XTR
MHQH29B-XTR
MHQH29B-XSR
MCX354A-QCAT
MCX353A-QCAT
NC543i (HP SL390 G7 in-built InfiniBand Host Channel Adapter)
CX-3 LOM down QDR
46M2199
46M2203



1.10 Software Supported

1.10.1 Remote Node Software Versions Supported in this Release

The Intel® True Scale Fabric Suite FastFabric management node can manage nodes with the following software:

- Host with FastFabric for OFED Enablement Tools 4.2 or later

Note:

While the Intel® True Scale Fabric Suite FastFabric Management Node requires Intel® OFED+ Host Software 1.5.3 or later to run Intel® True Scale Fabric Suite FastFabric 7.2, Intel® FastFabric can manage cluster nodes running Intel® OFED+ Host Software 1.2.5 or Intel® OFED+ Host Software 1.3, OFED 1.4 or OFED 1.5, and Intel® IB Tools 4.2 or later.

- Intel® Internally Managed 9000 series Switches with 4.1 or later firmware
- Intel® Externally Managed 9024FC Switches with 4.1 or later firmware
- Intel® Internally Managed 12000 series Switches with 5.0 or later firmware
- Intel® Externally Managed 12200 Switches with 5.0 or later firmware
- Intel® 12100 Switches with 5.0 or later firmware

The Intel® True Scale Fabric Suite Fabric Manager can manage nodes with the following software:

- Host with Intel® OFED+ Host Software 1.2 or later
- Intel® Internally Managed 9000 series Switches with 4.1 or later firmware
- Intel® Externally Managed 9024FC Switches with 4.1 or later firmware
- Intel® Internally Managed 12000 series Switches with 5.0 or later firmware
- Intel® Externally Managed 12200 Switches with 5.0 or later firmware
- Intel® 12100 Switches with 5.0 or later firmware


1.10.2 Remote Node Software Versions with Reduced Capability

The Intel® True Scale Fabric Suite FastFabric can manage nodes with the following software:

- Nodes running third-party IB Stacks
- OFED nodes without the Intel® IB Tools installed
- Third-party IB Switches

The Intel® True Scale Fabric Suite Fabric Manager can manage nodes with the following software:

- Nodes running third-party IB Stacks
- OFED Nodes with Intel® OFED+ 1.2 or earlier
- Third-party IB Switches



1.11 Installation Requirements

The following sections list any special or release-specific installation requirements for this release.

1.11.1 Package Installation Requirements:

Intel® True Scale Fabric Suite Software (IFS) package should be installed on the head node and Intel® OFED+ Host Software package should be installed on all other nodes except the head node.

When using Intel® True Scale Fabric Suite FastFabric toolset to install other nodes, `IntelIB-Basic.DISTRO.VERSION.tgz` should be downloaded. This file is specified by default in `fastfabric.conf` through the `FF_PRODUCT` and `FF_PRODUCT_VERSION` parameters and is used to install all other nodes.

1.11.2 Software and Firmware Requirements

All IFS software on a given node must be at the same release level. The Intel® OFED+ Host Software is installed as part of the package. Prior to installing the Intel® True Scale Fabric Suite Software release, any versions of the SilverStorm IB stack (and any other vendor's IB stack) must be uninstalled.

Note: When using the Intel® True Scale Fabric Suite (IFS) Software installation wrapper, the wrapper install enforces this requirement.

1.12 Changes for this Release

The following sections describe the changes that have been made to the Intel® True Scale Fabric Suite Software package since the last release.


For detailed information about any of the previous releases listed, refer to the Release Notes for the specific version.

1.12.1 Changes to Industry Standards Compliance

Table 1-7 shows each Basic OFED version that is supported and the Intel® OFED+ Releases that include each

Table 1-7. Changes to Industry Standards Compliance

Basic OFED Software Package Supported	Intel® OFED+ Host Software Package
Version 1.5.4.1	Versions 7.2.1.1.22, 7.2.2.0.8
Version 3.5-2	Version 7.3.0.0.26, 7.3.1.0.12
Version 3-12.1	Version 7.4.0.0.21
Version 3-18.1	Version 7.4.1.0.24
Version 3-18.2	Version 7.4.2.0.6
Version 3-18.3	Version 7.4.3.0.6
Version 4.8-1	Version 7.4.4.0.18



1.13 Product Constraints

The following is a list of product constraints for this release:

1.13.1 FastFabric Toolset Product Constraints

- The product supports a default HCAs configuration of Port 1 on the HCAs as the active port and Port 2 on the HCAs as the standby port. The following FastFabric operations may not work correctly with a HCA configuration of 2 active ports, or a configuration which has Port 2 of the HCAs as the active port:
 - Host Setup using FastFabric->Configure IPoIB IP Address
 - Host Admin using FastFabric->Verify Hosts ping via IPoIB
- All commands that are to be run on the chassis (Intel® and SilverStorm switches and gateways) should be invoked with the `-noprompt` option to avoid command execution time-out. This applies both to chassis commands invoked from the FastFabric TUI (Run a command on all chassis), as well as those invoked from the command line using the FastFabric `cmdall` command.

1.13.2 Fabric Manager

- Virtual Fabrics in this release leverage IBTA standard Partitioning Features. However, some OFED applications have limitations with regard to partitioning.
 - FastFabric – FastFabric tools are fully supported. Intel recommends that FastFabric be installed on an admin node which is a Member in the Default Partition (0xffff).
 - IPoIB – Intel recommends configuring Virtual Fabrics so that the first PKey on the port is the one desired for IPoIB on the host. Refer to the Configuration section of the *Intel® True Scale Fabric Software Installation Guide* for detailed information.
 - mvapich1 – To control the PKey, the `VIADDEV_DEFAULT_PKEY` must be exported at job startup. Refer to the Configuration section of the *Intel® True Scale Fabric Software Installation Guide* for detailed information.
 - Open MPI – To control the PKey, the `OMPI_MCA_btl_openib_ib_pkey` must be exported at job startup. Refer to the Configuration section of the *Intel® True Scale Fabric Software Installation Guide* for detailed information about this feature.
 - mvapich2 – To control the PKey, the `MV2_DEFAULT_PKEY` must be exported at job startup. Refer to the Configuration section of the *Intel® True Scale Fabric Software Installation Guide* for detailed information.

1.14 Product Limitations

There are no product limitations for this release.



1.15 Other Information

The following is a list of need-to-know information for this release:

- If FastFabric is being used, after an upgrade review the FF_PRODUCT parameter in `/etc/sysconfig/fastfabric.conf`. This parameter must be adjusted to match value shown in `/etc/sysconfig/fastfabric.conf-sample`.
- The `iba_rfm` command has been deprecated. For a new installation or an upgrade from an earlier release, instances of `iba_rfm` command will be removed. The `iba_top` command should now be used.

1.15.1 FastFabric Toolset Information

The FastFabric Toolset is automatically uninstalled if the base OFED release is uninstalled.

1.15.2 Fabric Manager Information

When there are many changes in the fabric (ISLs, switches going down) it is possible that many loops are no longer viable and the distribution of ISLs in the loops is becoming unbalanced. These changes can cause the loop test utilization to drop. Restarting of loop test will stop all traffic and compute fresh loop routes with balanced distribution of ISLs in loops.

1.15.3 Verbs Performance

Whenever possible, Intel recommends installing at least one True Scale HCA on a PCIe bus connected to the CPU 1 Socket for optimal verbs performance. For details on finding the correct PCIe slot, consult the documentation for your server chassis or motherboard.

1.16 Documentation

Table 1-8 lists the end-user documentation for the current release. All related documentation is available on the Intel download site.

Documentation for Intel® Partners is available at the vendors website.

Table 1-8. Related Documentation for this Release

Document Title
Intel® Hardware Documents
<i>Intel® True Scale 12000 Hardware Installation Guide</i>
<i>Intel® True Scale 12000 Users Guide</i>
<i>Intel® True Scale 12000 CLI Reference Guide</i>
<i>Intel® Adapter Hardware Installation Guide</i>
Intel® OFED+ Documents
<i>Intel® True Scale Fabric Software Installation Guide</i>
<i>Intel® OFED+ Host Software User Guide</i>
<i>Intel® OFED+ Host Software Release Notes</i>
Intel® IFS Documents
<i>Intel® True Scale Fabric Suite FastFabric User Guide</i>





Table 1-8. Related Documentation for this Release (Continued)

Document Title
<i>Intel® True Scale Fabric Suite Fabric Manager User Guide</i>
<i>Intel® True Scale Fabric Suite FastFabric Command Line Interface Reference Guide</i>
<i>Intel® True Scale Fabric Suite Software Release Notes</i>
Intel® Fabric Viewer Documents
<i>Intel® True Scale Fabric Suite Fabric Viewer Online Help</i>
<i>Intel® True Scale Fabric Suite Fabric Viewer Release Notes</i>



2.0 System Issues for Release 7.4.4.0.18

2.1 Introduction

This section provides a list of the resolved and open issues for this release of the True Scale Fabric Suite Software. The list of open issues includes a description and workaround.

2.2 Resolved Issues in this Release

Table 2-1 is a list of issues that are resolved in this release.

Table 2-1. Resolved Issues

Product/ Component/ PR Number	Release	Description
IFS/MPI 131369	7.4.4	For dual-rail systems, by default HCA2, port 1 is chosen for running traffic. If this port is disabled, mvapich2-2.1 fails to find any other active port. This causes mpirun to fail with error(s).
IFS 133007	7.4.4	Running the installation script with the parameter <code>--with-xeon-phi</code> provides XEON PHI support within the IFS installation process. The parameter <code>--with-xeon-phi=<path></code> accepts the path of where to look for <code>mpss-psm*</code> and <code>mpss-psm-dev*</code> packages which should be installed. If the parameter will be added without any path, the <code>- installer</code> script looks for <code>mpss-psm</code> packages within the default directories such as <code>/root</code> , <code>/tmp</code> , <code>/opt</code> or <code>/home</code> and automatically installs it.
True Scale/PSM 137021	7.4.4	When an app is linked to <code>libpsm_infinipath.so</code> , signals like <code>SIGSEGV</code> are not passed to the parent.
True Scale/IFS 134031	7.4.2	In order to realize good intra-node performance, the <code>kcopy</code> module must be installed. SLES blocks the module from being installed.



2.3 Known Issues

Table 2-2 lists the open issues for this release.

Table 2-2. Open Issues

Product/ Component/ PR Number	Description	Workaround
IFS/ Fabric Manager 111338	When the LogFile parameter is in use, the Fabric Manager outputs to the named file instead of syslog . If a high LogLevel is selected, the log file can grow quickly and consume too much disk space.	Limit use of LogFile to short duration debug type operations and use syslog for normal Fabric Manager operations.
IFS/ Rolls/Kits 123572	If the IFS kit is already installed, then running the updatenode command on the Installer and/or compute node (updatenode <headnode/compuetnode> command) returns errors similar to the following: <pre>compute000: Error: Package: opensm-devel-3.3.13-1.x86_64 (installed) compute000:Requires: opensm-libs = 3.3.13-1 compute000:Removing: opensm- libs-3.3.13-1.x86_64 (installed) compute000: opensm-libs = 3.3.13-1 compute000: Updated By: opensm- libs-3.3.15-1.el6.x86_64 (xCAT- rhels6.4-path0)</pre>	These errors may be safely ignored.
IFS 124980	If the libedit rpm is not present on a host, it should be installed prior to installing IFS. If the libedit libraries are not installed, the following error is displayed. Additionally, commands such as iba_manage_switch and iba_switch_admin would also fail with an error message similar to: <pre>#/opt/iba/ib_tools/L8simlnx /opt/iba/ib_tools/L8simlnx: error while loading shared libraries: libedit.so.0: cannot open shared object file: No such file or directory</pre> <p>If the libedit rpm is not installed prior to the IFS installation, the capture operation would fail. However, installing the libedit rpm would fix the issue. A reinstall of IFS is not needed.</p>	If not already present on the host, install the libedit rpm.
IFS/Other 133007	While running MPI between Xeon and Xeon-Phi within a host, the test fails with a scif_mmap_failed error.	Uninstall infinipath-libs and infinipath-devel rpms. Install mpss-psm and mpss-psm-dev rpms from mpss/psm directory.



Table 2-2. Open Issues (Continued)

Product/ Component/ PR Number	Description	Workaround
IFS/Install/ Uninstall 131977	<ul style="list-style-type: none"> • OFED-3.18-1 requires kernel-source package to be "installed" for SLES12SP1. • OFED-3.18-1 fails to start drivers if it does not find the kernel-source package during installation time. This is a requirement in order to use OFED-3.18-1 or IFS 7.4.1 on SLES12SP1. 	None
IFS/Install/ Uninstall 131984	If openibd is disabled using the iba_config (autostart none) option on SLES 12, after the system is rebooted systemctl should report openibd status as masked .	<p>The user would need to run iba_config option to re-enable it and reboot the system.</p> <p>If openibd is needed to be started manually, the user would need to execute unmask and enable the option systemctl to start openibd.</p> <p>Commands to manually start openibd would be</p> <pre> /usr/bin/systemctlunmask openibd /usr/bin/systemctl enable openibd /usr/bin/systemctlstart openibd </pre>
IFS/Install/ Uninstall 131780	Parallel studio 2016 is not supported in this release.	Continue using Parallel studio 2015.
IFS/Install/ Uninstall 131946	While installing IntelIBIFS.<osid>.7.4.1.0.* on RHEL and SLES systems, non-standard OFED RPMs may need to be removed before running the INSTALL script.	<p>Some examples are:</p> <ul style="list-style-type: none"> • qemu-kvm • gnome-boxes • libvirt-daemon-kvm • glusterfs-rdma • boost-devel • libboost_graph_parallel1_54_0 • libboost_mpi1_54_0 • libibnetdisc5
IFS/Install/ Uninstall 131005	SLES 12 systems have problems with modules not loading at boot.	On SLES12.x, set the allow_unsupported_modules parameter to 1 in the file /etc/modprobe.d/10-unsupportedmodules.conf .
True Scale/ Driver 131486	rdma-ndd is a background service; it updates the node description as hostname devicename . However, the service openibd already sets the appropriate node description as Hostname HCA_n .	Since rdma-ndd overwrites the correct node description, it needs to be disabled in order to expect correct behavior by the openibd service.



Table 2-2. Open Issues (Continued)

Product/ Component/ PR Number	Description	Workaround
IFS/ FastFabric 131984	For a SLES12 host, if <code>openibd</code> is disabled using <code>iba_config</code> , (<code>autostart none</code>) option, after the system is rebooted <code>systemctl</code> would report <code>openibd</code> status as masked: <pre>CMZ-DEV-114:~ # systemctl status openibd openibd.service Loaded: masked (/dev/null) Active: inactive (dead)</pre>	To re-enable it and reboot the system, run the <code>iba_config</code> option. If <code>openibd</code> needs to be started manually, the user needs to execute the <code>unmask</code> and <code>enable</code> options of <code>systemctl</code> to start <code>openibd</code> . Commands to manually start <code>openibd</code> would be / <pre>usr/bin/systemctl unmask openibd /usr/bin/systemctl enable openibd /usr/bin/systemctl start openibd</pre>
True Scale/ MPI 135513	Mvapich2 compiled with <code>icc</code> or <code>pgi</code> can fail when using <code>mpirun_rsh</code> .	Use the SSH-based version - <code>mpirun</code> .
True Scale/ Driver 135817	Because of a huge I/O mapping within the kernel, MPSS on RHEL7.3 (and above) and SLES 12 SP2 crashes.	Using <code>nohugeiomap</code> in the kernel boot parameters allows MPSS to function properly.
Third Party/ RHEL/OFED 139608 141418	The OFED IPoIB driver from the 4.8 kernel is not fully-compatible with the bonding drivers distributed in SLES 12.1 and 12.2 (causing kernel-panic), as well as RHEL 7.4 (bonding fails to activate).	None.

